

Analyzing Evolving Stories in News Articles

Roberto Camacho Barranco
The University of Texas at El Paso
500 W University Ave.
El Paso, Texas 79968
rcamachobarranco@utep.edu

Arnold P. Boedihardjo
U. S. Army Corps of Engineers
7701 Telegraph Rd
Alexandria, Virginia 22315
arnold.p.boedihardjo@usace.army.mil

M. Shahriar Hossain
The University of Texas at El Paso
500 W University Ave.
El Paso, Texas 79968
mhossain@utep.edu

ABSTRACT

There is an overwhelming number of news articles published every day around the globe. Following the evolution of a news-story is a difficult task given that there is no such mechanism available to track back in time to study the diffusion of the relevant events in digital news feeds. The techniques developed so far to extract meaningful information from a massive corpus rely on similarity search, which results in a myopic loopback to the same topic without providing the needed insights to hypothesize the origin of a story that may be completely different than the news today. In this paper, we present an algorithm that mines historical data to detect the origin of an event, segments the timeline into disjoint groups of coherent news articles, and outlines the most important documents in a timeline with a soft probability to provide a better understanding of the evolution of a story. Qualitative and quantitative approaches to evaluate our framework demonstrate that our algorithm discovers statistically significant and meaningful stories in reasonable time. Additionally, a relevant case study on a set of news articles demonstrates that the generated output of the algorithm holds the promise to aid prediction of future entities in a story.

CCS CONCEPTS

•Information systems → Data mining; Summarization; Document topic models; •Mathematics of computing → Continuous optimization;

KEYWORDS

Entity evolution, connecting the dots, storytelling, entity prediction, document understanding, causal analysis

ACM Reference format:

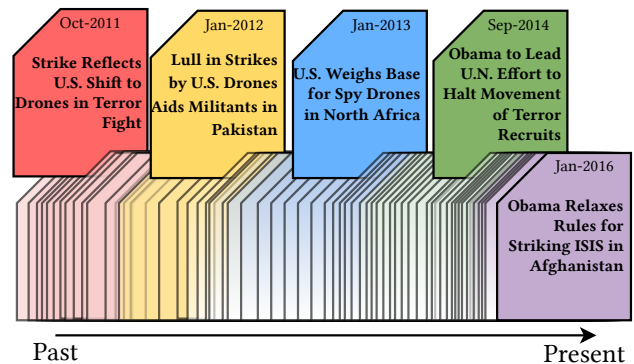
Roberto Camacho Barranco, Arnold P. Boedihardjo, and M. Shahriar Hossain. 2016. Analyzing Evolving Stories in News Articles. In *Proceedings of ACM Conference, Washington, DC, USA, July 2017 (Conference'17)*, 9 pages. DOI: 10.475/123_4

1 INTRODUCTION

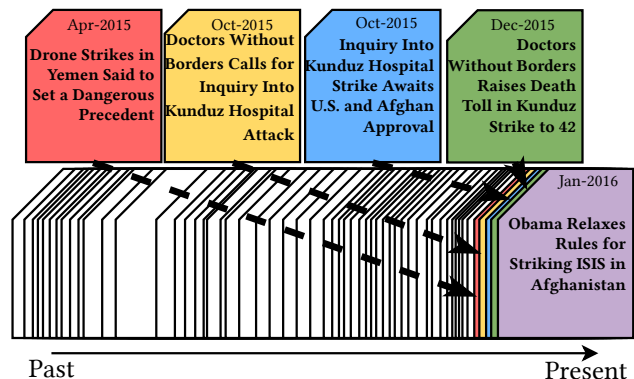
The pervasiveness of the Internet has greatly facilitated editors to publish a large number of news articles everyday in any topic,

practically without any restrictions of page-limits. Thus, the number of articles available to avid readers increases every day at a very fast pace. Many business and government organizations track these news articles to study public sentiments, business directions, progression of events, and many other aspects of economic and socio-political issues. However, tracking topics from different perspectives is a difficult task given that one news-story of today's interest could have evolved from another story of the past, thus forming a chain of relevant events. This leads to the concept of *diffusion theory* [4] which refers to the change of the distributional patterns of a phenomenon over time.

While capturing such diffusions over a timeline is still a challenge, researchers have targeted the problem of tracking stories in different guises, e.g., storytelling [14, 20], storyboarding [23], connecting



(a) A diffusion-based story.



(b) A similarity-based story.

Figure 1: A diffusion-based approach captures documents from the past reflecting a smooth transition of the evolving story. In contrast, a similarity-based approach focuses on a narrow range of the timeline.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Conference'17, Washington, DC, USA

© 2016 Copyright held by the owner/author(s). 123-4567-24-567/08/06...\$15.00
DOI: 10.475/123_4

the dots [15, 30], and metro maps [31]. Most of these methods find underlying connections between articles using a similarity-based network of documents. Their objective is mainly focused on generating a cohesive thread of a story. An extensive usage of cohesion may result in stories that revolve around a single *chapter* of the evolving thread of the story. Figure 1 (a) illustrates a scenario generated using our diffusion-based approach, which captures the articles published in the past reflecting the evolution of the story. Figure 1(b) presents the results obtained, for the same seed document, using a similarity-based approach to track past documents similar to the article of interest. While our approach brings out the underlying diffusion of concepts and their progression over time, the similarity-based approach provides a coherent story that revolves around a single event — a drone strike in Kunduz, as shown in both examples.

This paper presents a novel framework that takes as input a large corpus of documents, along with a user-specified set of seed documents. The framework runs through a palette of natural language processing (NLP) and other preprocessing tasks before it reaches its core objective function, which generates a storyline by finding and combining turning points that lead to the seed story. The end-product of the framework is a concise chain of documents from a multitude of news articles to help the user identify the evolution of a story, as shown in Fig. 1(a).

The contributions of this paper can be summarized as follows:

- **Evolution:** We propose a novel method to discover a chain of documents published in the past given a set of seed documents, where the chain reflects evolution of the concepts in the seed documents.
- **Diffusion and similarity:** Our proposed technique captures the underlying diffusion of concepts and has the flexibility to incorporate the similarity concepts from the state-of-the-art methods.
- **Prediction:** We demonstrate the potential of the outcome generated by our approach as a tool to predict future states of an evolving story.
- **Evaluation:** We conduct a set of experiments to quantitatively evaluate different aspects of our model. Additionally, a human participant-based study demonstrates that the discovered story lines were satisfactory.

2 RELATED WORK

Several problems and tasks related to our work have been well studied. These problems include *summarization*, *event detection*, *event extraction*, *event evolution*, *event prediction* and *storytelling*, also known as *connecting the dots*. *Summarization* consists of extracting the most important sentences from a text corpus to help the user understand the main idea of the set of documents [3, 5, 25, 27, 35].

Event detection consists of detecting temporal bursts of highly correlated documents [18, 36, 37], while *event extraction* aims at obtaining useful knowledge regarding these events [21, 29]. *Event evolution* [34, 38] and *topic evolution* [16, 17, 33] study how these abstract concepts (*events* and *topics*) evolve over time, based on latent relations within the corpus and temporal information encoded in each document.

The *event prediction* through text mining problem has also been addressed in the literature. Radinsky *et al.* [28] described a system that uses *causality extraction* to obtain pairs of terms that have a causal relation and that are later used to train a prediction algorithm. Luo *et al.* [22] also address this problem by introducing the concept of *semantic uncertainty*, which is used to estimate the *most certain* next state based on the current state, or event. This approach is particularly useful when there is limited historical data available.

The main component of our framework is focused on the problem of *storytelling*, or *connecting the dots*. There are several approaches that aim to solve this problem for different types of datasets, such as scientific articles [15], entity networks [7, 8, 13], image collections [12, 32] and document collections [1, 10, 11, 20, 24, 30, 31, 34, 40]. Most of the work on this problem uses graph-based representations of a document or entity set [7, 10, 14, 24, 40]. *Storytelling* has several different applications, such as intelligence analysis [13, 14], news recommendation [11], search engines [8, 40], and social network analysis [7, 26, 32].

Solutions to the *connecting the dots* problem include the use of probabilistic approaches, such as random walks [11, 30, 31, 34, 40], Monte Carlo simulations [1] and determinantal point processes (DPP) [10]. Shahaf *et al.* [30] introduced the concept of coherence and coverage to assess the quality of a chain limited by two user-specified endpoints, and then extended this work by building *metro maps* [31] which are formed by several coherent chains that intersect at some points, forming a map-like structure. Gillenwater *et al.* also obtain several coherent chains using a DPP-based model.

Our work has several differences with respect to the previously mentioned methods. First, we use an optimization-based approach similar to the one presented by Shahaf *et al.* [30], but we introduce the concept of *diffusion* to identify the evolution of the concepts in the seed documents. The use of diffusion allows our framework to create stories based on a combination of different events, which is not possible using previously proposed methods. Second, our work discovers the most relevant documents associated with each turning point to provide a concise summary for each turning point event. Third, we leverage the topic distribution of each document to reduce our search space while keeping important documents for our optimization. Finally, we suggest the possibility of using the resulting relevant documents to build a regression framework for entity prediction.

3 PROBLEM DESCRIPTION

This work focuses on news articles and the entities within the news. The entities detected are persons, organizations, and locations. Let $\mathcal{D} = \{d_1, d_2, \dots, d_{|\mathcal{D}|}\}$ be the set of documents and $\mathcal{E} = \{e_1, e_2, \dots, e_{|\mathcal{E}|}\}$ be the set of entities in the news corpus. Each document $d \in \mathcal{D}$ contains a set of entities $\mathcal{E}_d \subset \mathcal{E}$. Additionally, each document d contains its publication date t_d . As part of the preprocessing steps, we apply topic modeling using Latent Dirichlet Allocation (LDA) [6] on \mathcal{D} , obtaining a topic distribution T_d for each document d .

3.1 What is the expected outcome?

Let the user input be a set of seed documents $\mathcal{D}_{seed} \subset \mathcal{D}$, where $|\mathcal{D}_{seed}| \geq 1$. We define a *turning point* in the story $\tau \in \mathcal{T}$ as a

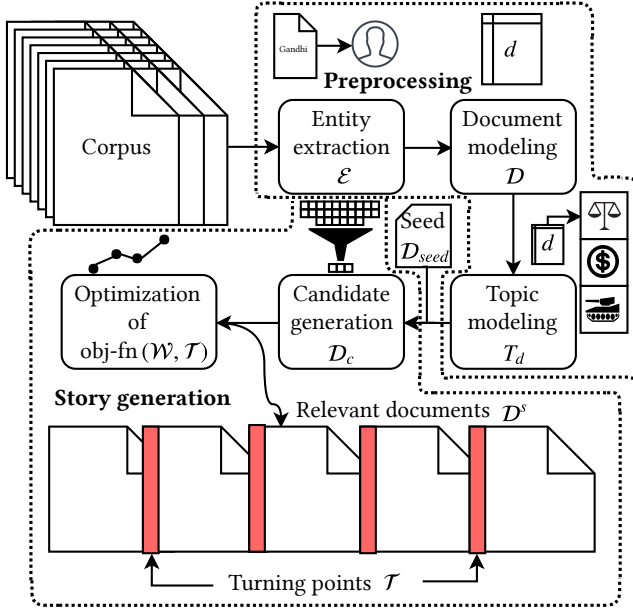


Figure 2: Proposed framework.

specific date in which the story under analysis has a significant change. Let S be a vector where each element, defined as *segment*, is a pair of consecutive turning points such that $|S| = |\mathcal{T}| - 1$, and $S = [(t_1, t_2), (t_2, t_3), \dots, (t_{|\mathcal{T}|-1}, t_{|\mathcal{T}|})]$.

Our main goal is to split the set of documents in $|S|$ segments by finding $|\mathcal{T}| - 2$ turning points, since the first and last turning points are fixed. For each segment $s \in S$ we want to find a subset of documents $D^s \subset D$ that help the user study the evolution of the prominent entities in D_{seed} through various different events.

4 METHODOLOGY

Our framework comprises two main stages: (1) preprocessing, which creates document and topic models from a text corpus and (2) story generation, which takes a set of seed documents and other constraints as inputs and, via an optimization routine, outputs a story formed by relevant documents that have a common thread, but that belong to different events. Fig. 2 illustrates this framework.

4.1 Preprocessing

In the preprocessing stage, our framework: (1) extracts entities (e.g., person, location, organization) from the documents, (2) represents each document as a vector of entities, and (3) obtains a topic distribution for each document in our corpus. We extract entities from the text corpus using standard Named Entity Recognizers [2, 9]. Our framework leverages a tf-idf model with cosine normalization [14] to generate weights $w(e, d)$ for each entity $e \in \mathcal{E}$ in each document $d \in D$.

4.2 Story generation

In this stage, our framework furnishes a candidate set of documents $D_c \subset D$ from the entire dataset that satisfy some temporal and topical constraints with respect to the seed documents D_{seed} . The temporal and topical constraints are driven by user input regarding how far back in time the algorithm should track to detect an origin,

or, how much deviation the algorithm should allow in terms of topic distribution as compared with the seed set.

The temporal criteria establishes that all of the candidate documents must have been published before the most recent seed document and within a certain minimum threshold t_{\min} , i.e., d_i can only be part of D_c if $t_{\min} < t_i < \max(t_{seed})$, where t_i is the publication date of document d_i . The topical criteria expresses, in terms of topic distributions, how much deviation the candidate documents can have from the seed documents when optimizing for a diffusion and cohesion objective. For each document $d_{seed} \in D_{seed}$, we compute the KL-divergence [19], or *topical divergence*, between the topic distribution of the seed document $T_{d_{seed}}$ and the T_d of all the documents $d \in D$. Document d is included in D_c , only if $\forall d_{seed} \in D_{seed} \left(\text{KL-divergence}(T_d, T_{d_{seed}}) \leq \alpha \right)$, where T_d is the topic distribution of document D , obtained using the LDA [6] and α is a user-defined parameter. In Section 5, we will refer to D_c by D , for simplicity.

Finally, the framework optimizes fitment of a story model that identifies turning points of the story over a timeline and provides relevant documents within each segment. The outcome of the optimization routine helps understand the evolution of the news-story described in the seed document(s). This essential is further described in Section 5.

5 FORMALIZING OUR APPROACH

In this section, we gradually introduce the concepts and formulations that give shape to our model, which has the goal of obtaining a coherent story from separate events. We present our ideas iteratively and explain why and how new penalty terms were developed to achieve our goal.

5.1 Defining incoherence

The objective is to ensure that the documents within each time segment of S are coherent. A simplistic approach is to minimize the distance of all pairs of documents within each time segment. For simplicity, we use the notation d^s , to indicate the segment s to which document d belongs to, i.e., d^1 is defined as a document in the segment between τ_1 and τ_2 . We define this distance as incoherence:

$$\text{incoherence}_1(s) = \frac{\sum_{i \neq j}^{|P^s|} \|d_i^s - d_j^s\|_{\varrho}}{|P^s|} \quad (1)$$

where S is the set of segments, d^s is a document s.t. $d^s \in D^s$, where D^s is the set of candidate documents that are within segment s . P^s is a set of pairs (i, j) that represent the different 2-permutations of the indices of pairs of documents d_i^s and d_j^s , computed using the following equation.

$$|P^s| = \frac{|D^s|!}{2(|D^s|-2)!} = \frac{|D^s|^2 - |D^s|}{2} \quad (2)$$

5.2 Defining a distance metric and forming an initial objective function

In Eq. 1, the distance computation between the weight vectors of two documents is represented as $\|\cdot\|_{\varrho}$, where ϱ is substituted by the distance measure chosen. We prefer the Soergel distance between a pair of documents d_i and d_j , defined in Eq. 3.

$$\text{soergel}(d_i, d_j) = \frac{\sum_{e \in \mathcal{E}} |w(d_i, t) - w(d_j, t)|}{\sum_{e \in \mathcal{E}} \max(w(d_i, t), w(d_j, t))} \quad (3)$$

where \mathcal{E} is the set of all the entities that appear in the dataset, and $w(d, e)$ is the tf-idf weight of entity e for document d .

The initial objective function is thus defined as Eq. 4, where s represents the current segment.

$$\mathcal{F}_1(\mathcal{T}) = \sum_{s=1}^{|S|} \text{incoherence}_1(s) \quad (4)$$

This initial objective function is not enough to serve the purpose of capturing evolution in terms of diffusion. Eq. 4 tends to generate one very large segment and many small ones, which is undesirable. Ideally, two documents that were published at distant dates should not belong to the same segment. Another limitation of the approach in Eq. 4 is that it uses hard assignments of documents to a single segment, while we would prefer a soft membership of documents to time segments for better analytic flexibility.

5.3 Including penalty based on time segmentation

To tackle the problem of generating small segments, we include a new term in our definition of *incoherence*. The new term helps by avoiding segments containing document pairs with distant publication dates. Therefore, a solution with many small segments or a solution with a very large segment will not be the optimal. This term is simply a date difference penalty (date_Δ).

$$\text{date}_\Delta(i, j) = |t_i - t_j| \quad (5)$$

where t_i is the publication date of document d_i . Although Eq. 5 provides date differences, in practice, we normalize the values to a smaller range and the user can flexibly modify the normalization range. With this new term, we can redefine incoherence as Eq. 6.

$$\text{incoherence}_2(s) = \frac{\sum_{i < j}^{|P^s|} \text{soergel}(d_i, d_j) * \text{date}_\Delta(i, j)}{|P^s|} \quad (6)$$

The objective function has the same form of Eq. 4, but uses the new incoherence equation. The advantage of this new formulation is that it reduces the incidence of very small segments significantly. However, the current objective function is not capable of capturing the essence of *diffusion* or *turning point* because it only takes into account coherence of each segment. This can result in very similar segments that do not show the evolution of a story. Another disadvantage is that a document can only be assigned to a single segment, and thus only discrete optimization routines can be used.

5.4 Introducing unconnectedness

While we prefer similar documents within a segment, we expect high dissimilarity between the documents of different segments. We call this property of diffusion across segments *unconnectedness*. We expect high *unconnectedness* as expressed by Eq. 7.

$$\text{unconnectedness}(s) = \frac{\sum_{i=1}^{|D^s|} \sum_{j=1}^{|G^s|} \text{soergel}(d_i^s, g_j^s)}{|D^s| * |G^s|} \quad (7)$$

where $G^s = D - D^s$ is the set of documents that are not within segment s , and $g^s \in G^s$. Combining unconnectedness with the original objective function, we obtain the following equation.

$$\mathcal{F}_2(\mathcal{T}) = \sum_{s=1}^{|S|} (\text{incoherence}_2(s) - \text{unconnectedness}(s)) \quad (8)$$

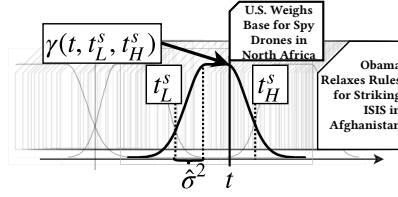


Figure 3: The $\gamma(t, t_L^s, t_H^s)$ function as described in Eq. 11.

In practice, we normalized and rescaled these two terms to adjust the importance of minimizing incoherence and maximizing unconnectedness. However, a major problem with this formulation is that because of the conflicting objectives, we cannot easily define which penalty should have the largest effect.

5.5 Changing to segment similarity

To eliminate the conflicting objectives issue, and to obtain a smooth formulation, we introduce the concept of similarity between segments. This penalty, formalized as Eq. 9, will measure how similar documents from one segment are with respect to the documents in the other segments.

$$\text{similarity}_1(s) = \frac{\sum_{i=1}^{|D^s|} \sum_{j=1}^{|G^s|} e^{-\text{soergel}(d_i^s, g_j^s)}}{|D^s| * |G^s|} \quad (9)$$

We seek to minimize similarity(s) for each segment s . The updated objective function is:

$$\mathcal{F}_3(\mathcal{T}) = \sum_{s=1}^{|S|} (\text{incoherence}_2(s) * \text{similarity}_1(s)) \quad (10)$$

Eq. 10 allows minimization of incoherence and similarity simultaneously without keeping track of the importance of each other since they are represented as factors. However, this formulation still suffers from two issues: (1) a news article can only belong to a single segment, which means only discrete optimization routines can be used, and (2) it is possible to have several articles that are irrelevant to the story within each segment. We address these issues in the following subsections.

5.6 From discrete to continuous

The objective function of Eq. 10 considers discrete assignments of documents to time segments. This results in lesser flexibility in terms of smoothness and use of optimization routines. To indicate the certainty that a document with time stamp t belongs to a continuous range of time, in our case a segment, where t_L^s is the lower limit or a turning point and t_H^s is the upper limit, a membership function can be defined as:

$$\gamma(t, t_L^s, t_H^s) = \begin{cases} \frac{1}{\sqrt{2\pi\hat{\sigma}^2}} e^{-\frac{(t-t_L^s-\hat{\sigma}^2)^2}{2\hat{\sigma}^2}} & \text{if } t \leq t_L^s \\ \frac{1}{\sqrt{2\pi\hat{\sigma}^2}} & \text{if } t_L^s < t < t_H^s \\ \frac{1}{\sqrt{2\pi\hat{\sigma}^2}} e^{-\frac{(t-t_H^s+\hat{\sigma}^2)^2}{2\hat{\sigma}^2}} & \text{if } t_H^s \leq t \end{cases} \quad (11)$$

Although Eq. 11 is partitioned, it is continuous and smooth because the membership function is formed by two halves of different Gaussian distributions, and a linear section that connects both, as shown in Fig. 3. The means of the distributions are the lower and

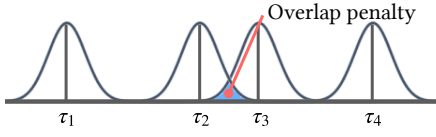


Figure 4: Potential overlap of turning points.

higher limits, and the standard deviation $\hat{\sigma}$ for both is selected by the user to modify the shape of the membership function. This score is not a probability, but we can obtain the probability that time stamp t belongs to segment s by dividing the score over the sum of the membership values of t for all segments in \mathcal{S} .

With this new concept of membership function, we must update our previous definitions.

$$\text{incoherence}_3(s) = \frac{\sum_{i,j}^{|D| \times |D|} \Phi * \text{soergel}(d_i, d_j) * \text{date}_\Delta(i, j)}{\sum_{i,j}^{|D| \times |D|} \Phi} \quad (12)$$

where

$$\Phi = \gamma(t_i, t_L^s, t_H^s) * \gamma(t_j, t_L^s, t_H^s) \quad (13)$$

Segment similarity can be redefined as:

$$\text{similarity}_2(s) = \frac{\sum_{i,j}^{|D| \times |D|} \phi * e^{-\text{soergel}(d_i, d_j)}}{\sum_{i,j}^{|D| \times |D|} \phi} \quad (14)$$

where

$$\phi = \gamma(t_i, t_L^s, t_H^s) * (1 - \gamma(t_j, t_L^s, t_H^s)) \quad (15)$$

The objective function has the same form of Eq. 10, but uses the new incoherence and similarity equations. The main advantage of this new formulation is that it is suitable to be used with a larger variety of optimization routines. However, a new problem is introduced with this objective function, which is that it tends to overlap the turning points. Furthermore, we still have the problem of many documents being irrelevant to a segment.

5.7 Introducing segment overlap penalty

To avoid overlap of segments (as shown in Fig. 4) we include a penalty that is high when two turning points are very close to each other. It is defined as:

$$\text{overlap} = \left(1 + \sum_{m,n,m < n}^{|S|-1 \times |S|-1} e^{-\left(\frac{(t^m - t^n)^2}{2\sigma^2} \right)} \right) \quad (16)$$

where t^m and t^n are the publication date of two different documents, and σ is a hyperparameter that allows the user to set a threshold defining the degree of the allowed overlap.

The objective function with the overlap penalty is:

$$\mathcal{F}_4(\mathcal{T}) = \sum_{t=1}^{|S|} (\text{incoherence}_3(s) * \text{similarity}_2(s)) * \text{overlap} \quad (17)$$

5.8 Adding weights as a relevance factor

It is difficult for a user to analyze a massive amount of candidate documents given only the time segments. Our framework mitigates this problem by providing a score with each document representing how important it is for this study. This score is a variable for the optimization routine and it is used in the objective function. The variable is represented as a vector \mathcal{W} of length equal to the number of candidate documents. A weight of 1.0 indicates that the document is relevant and 0.0 indicates that it is not.

The updated definitions are Eq. 18 and Eq. 19, where $w_i \in \mathcal{W}$ is the weight for document d_i .

$$\text{incoherence}_4(s) = \frac{\sum_{i,j}^{|D| \times |D|} w_i * w_j * \Phi * \text{soergel}(d_i, d_j) * \text{date}_\Delta(i, j)}{\sum_{i,j}^{|D| \times |D|} w_i * w_j * \Phi} \quad (18)$$

$$\text{similarity}_3(s) = \frac{\sum_{i,j}^{|D| \times |D|} w_i * w_j * \phi * e^{-\text{soergel}(d_i, d_j)}}{\sum_{i,j}^{|D| \times |D|} w_i * w_j * \phi} \quad (19)$$

The objective function has the same form of Eq. 17, but uses the new incoherence and similarity formulations and has \mathcal{W} as input. However, a new problem emerges — the optimal value is when all the weights are set to zero.

5.9 Uniformity penalty

To ensure that only a subset of \mathcal{D} is truly relevant, and to avoid the degenerate case where the optimal value is with all the weights set to zero, we need to minimize the uniformity of the weights vector \mathcal{W} . A completely uniform vector \mathcal{W} , i.e., all zeros or all ones, will result in a high penalty, while a completely non-uniform weights vector will result in no penalty. The uniformity penalty is:

$$\text{uniformity} = \left(1 + \sum_{s=1}^{|S|} \left(1 - \frac{\left(\left\| \frac{\mathcal{W}_s * \Gamma_s^\top}{\sum \mathcal{W}_s * \Gamma_s^\top} \right\|_2 * \sqrt{|\mathcal{W}_s|} \right) - 1}{\sqrt{|\mathcal{W}_s|} - 1} \right) \right) \quad (20)$$

where Γ_e is a vector of values returned by the membership function (Eq. 11 for the documents that fall in segment s).

5.10 The final objective function

The final objective function is shown below.

$$\mathcal{F}_5(\mathcal{T}, \mathcal{W}) = \sum_{s=1}^{|S|} (\text{incoherence}_4(s) * \text{similarity}_3(s)) * \text{overlap} * \text{uniformity} \quad (21)$$

The objective function is minimized for two vectors: \mathcal{S} and \mathcal{W} . The elements $s \in \mathcal{S}$ are bounded between $[0, \text{date}_{\max}]$ where date_{\max} is the value representing the publication date of the most recent news article while 0 represents the earliest document on the candidate dataset. The elements $w \in \mathcal{W}$ are bounded between $[0, 1]$, as stated earlier. We used the quasi-newton limited memory algorithm for bound constrained optimization (L-BFGS-B) [39] to minimize our objective function.

6 EXPERIMENTAL RESULTS

For our experiments, we use a corpus of 400,842 New York Times articles published between January 2000 and June 2016. The corpus contains 3,320,886 unique entities. To the best of our knowledge, there are no publicly available labeled benchmark datasets that can allow us to perform a supervised evaluation of our framework. Our evaluations are based on a combination of statistics- and human participant-based studies. We seek to answer the following questions.

- (1) How well does our approach reflect the evolution of a story? (Section 6.1)
- (2) How does our method compare to other methods such as clustering and similarity-based storytelling? (Section 6.2)
- (3) How do the various proposed regularizations impact story quality? (Section 6.3)

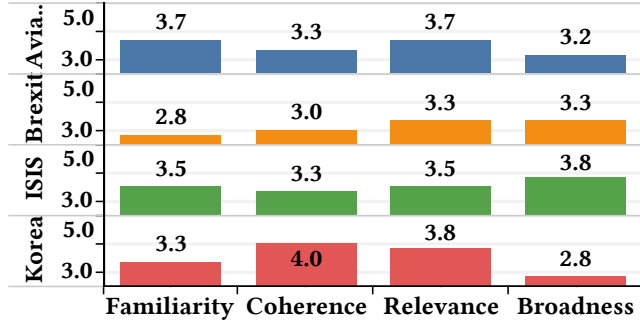


Figure 5: Average scores obtained for each metric per topic, a higher score is better (max. 5.0).

- (4) How does the statistical significance for the position of the turning points change while modifying the different hyper-parameters of our optimization? (Section 6.4)
- (5) What is the repeatability of the optimization? (Section 6.5)
- (6) Can we use the set of highly relevant documents with respect to a seed document to predict which entities are expected to appear in a future document? (Section 6.6)

6.1 Evaluation of Chain Continuity

The lack of a benchmark dataset makes it very difficult to evaluate the performance of our optimization framework. To address this issue, we designed a user study motivated by the work of Shahaf *et al.* [30]. We picked four different topics and two different random documents related to each topic as seed documents. The topics selected were *aviation*, *Brexit*, *ISIS*, and *tensions between North and South Korea*. Four evaluators were asked to score three topics each, so that each topic had a total of six evaluations¹.

For each segment, we selected the document with the highest weight. Four past documents were selected for each seed document. Users evaluated four criteria on a scale of 1 to 5, with five being the best, for each sequence:

- *Familiarity*: How familiar the person is with this topic.
- *Relevance*: Do the documents seem relevant to the topic?
- *Coherence*: Is the chain coherent? A lower score is provided if there is a document in the chain that does not belong to the topic.
- *Broadness*: Does each document provide new information guaranteeing a slight shift in topic over time?

The average scores per criterion are presented in Fig. 5. The results indicate that our algorithm performs, on average, above a score of 3, which indicates that the chains are coherent, relevant, and broad enough to capture evolution. An interesting observation is that the *Brexit* topic, which had the lowest average familiarity score, also has the lowest scores for coherence and relevance, indicating that such a study might be biased toward the familiarity of the topic.

6.2 Comparison with Other Methods

We compare the outcomes of our framework with two other methods: (1) a *k*-means clustering, which combines the time feature with entities and (2) a similarity-based approach where consecutive

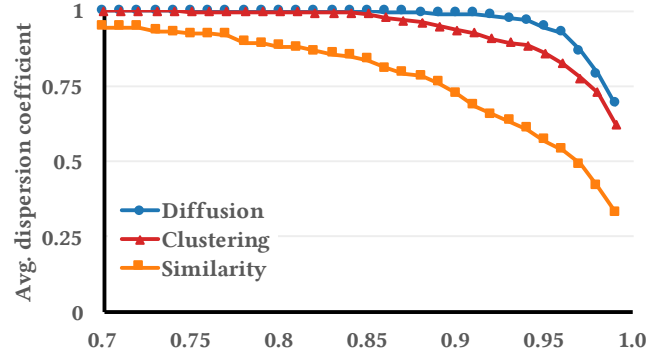


Figure 6: Comparison of average dispersion coefficient for stories created using our diffusion based approach, a clustering based algorithm, and a similarity based technique. Our approach results in the best (highest) dispersion.

nearest neighbors are selected as a chain starting with the seed document. We selected one best document per time-segment for this evaluation using our method. For the *k*-means clustering based approach we selected one document closest to the centroid of each cluster as a relevant document. For the similarity-based approach, we used the sequence of documents generated from the seed document as a path. Notice that the two alternative approaches focus on distance or similarity, while our objective function is designed to capture diffusion. Hossain *et al.* introduced the concept of *dispersion coefficient* in evaluating the Storytelling algorithm [14]. We use the same concept to evaluate the quality of a chain of documents $\{d_0, d_1, \dots, d_{n-1}\}$, containing n articles.

$$\psi = 1 - \frac{1}{n-2} \sum_{i=0}^{n-3} \sum_{j=i+2}^{n-1} \text{disp}(d_i, d_j)$$

where

$$\text{disp}(d_i, d_j) = \begin{cases} \frac{1}{n+i-j} & \text{if } \text{soergel}(d_i, d_j) < \theta \\ 0 & \text{otherwise} \end{cases}$$

The dispersion coefficient of a chain of documents is 1.0 only if consecutive pairs of documents meet a distance threshold θ . The coefficient is 0 when every pair of documents in the chain satisfies the distance threshold.

We generated chains for 116 different seed documents using all three approaches under consideration. Figure 6 compares the average dispersion coefficient versus the distance threshold. Figure 6 shows that the average dispersion of our diffusion-based method provides the highest dispersion coefficients for any distance threshold. This indicates that our method generates chains with smooth transition of topics, which is our goal.

6.3 Use Case Study

In this experiment, we perform a use case study to compare the results of our diffusion-based approach with two preliminary versions of our objective function, as well as with the similarity-based approach described in Sec. 6.2. We execute each of these methods for the same seed document and select the document in the middle of each segment. The number of segments is set to 5. Fig. 7(a) shows the story generated using the formulation presented in Sec. 5.2. In this case, because the date penalty was not included, one event overlaps with the other, and thus only four segments are obtained.

¹Evaluation questions available at: <https://storyeval.herokuapp.com/>

(a)		Nov-2011 Euro Zone Leaders Weigh New Budget Rules	Jul-2013 E.U. Leaders Discuss Ways to Tackle Youth Unemployment	Dec-2015 E.U. Agrees to Share Air Passenger Data	May-2016 As Brexit Vote Nears, British Data Shows Increase in Immigration
(b)	Apr-2012 The Netherlands: Call for Elections	Sep-2013 Re-election in Hand, Merkel Moves to Form New German Government	Jan-2014 Europe, Facing Economic Pain, May Ease Climate Rules	Dec-2014 That Debt From 1720? Britain's Payment Is Coming	May-2016 As Brexit Vote Nears, British Data Shows Increase in Immigration
(c)	Dec-2011 Leaders Look to I.M.F., Again, as Euro Crisis Lingers	May-2012 Rising Greek Political Star, Foe of Austerity, Puts Europe on Edge	Sep-2013 Re-election in Hand, Merkel Moves to Form New German Government	Oct-2014 Bloc in Europe Starts to Balk Over Austerity	May-2016 As Brexit Vote Nears, British Data Shows Increase in Immigration
(d)	Feb-2016 What's at Stake as Britain and E.U. Discuss How to Maintain a Union	Apr-2016 Uncertain Economics Influence Brexit Talk	Apr-2016 Britons Will Be Poorer if They Leave E.U., Government Asserts	May-2016 David Cameron Says Leaving E.U. Would Hobble British Economy	May-2016 As Brexit Vote Nears, British Data Shows Increase in Immigration

Figure 7: Stories generated using (a) the objective function presented in Sec. 5.2, (b) the formulation presented in Sec. 5.5, (c) our diffusion-based approach and (d) a similarity-based approach.

The topics of each document are about the European Union, but lack a common thread. Fig. 7(b) shows a story generated using by optimizing the objective function presented in Sec. 5.5, which is in the continuous space and includes the date and similarity penalties. An analysis of this story shows that there is no coherent thread, which is caused by the lack of weights that indicate the relevance of a document.

Fig. 7(c) shows the story resulting from our diffusion-based approach. At first glance, some of the articles seem disconnected from a continuous thread. However, after a brief analysis, we can find a very relevant thread: the first article is about the financial assistance required by Greece and Portugal because of the severe crisis. The second article mentions how Germany made efforts to keep Greece in the Euro zone, because of fear of the consequences of Greece leaving. The next article mentions the future plans of the German chancellor with respect to the European Union and the imposed austerity, and in particular mentions that Britain's prime minister expected the support from Ms. Merkel. The article before the seed document mentions the latest European crisis, and how several countries were "rebellious" against the German-mandated austerity. The seed document talks about immigration being one of the main motivations for Brexit. The article mentions that there was an influx of immigrants on 2015, presumably escaping from the economical crisis in their countries of origin. This is a coherent and relevant thread consisting of different events, which was our ultimate goal.

The story obtained using the similarity-based approach, shown in Fig. 7(d), is also coherent and relevant. However, the story completely revolves around Brexit, and a simple keyword search about the topic would probably return similar results. Thus, we consider that our result provides a better insight into the possible paths that lead to the event described in the seed document.

6.4 Statistical Significance Analysis

Evaluation using precision, recall, and other variants does not apply in the context of our work because of the lack of labeled datasets suitable for our task. An alternative approach, other than the user studies and quality-evaluation using an unsupervised metric like dispersion coefficient, is to use a statistical analysis of the significance (p-value) as a sanity check to make sure that our formulation obtains results that have a very low probability of being obtained by random chance. Our objective function has several user-modifiable parameters that can change analytic viewpoints. In this section, we will explore how the significance changes with varying parameters.

For our problem, we defined p_{TP} as the p-value for the turning points vector \mathcal{T} . We generated m random samples \mathcal{T}_s which are of the same size as \mathcal{T} . Then, we compared each of the m random samples with \mathcal{T} by computing their element-wise difference and comparing it to a tolerance β , i.e., $(\mathcal{T}_s - \mathcal{T}) \leq \beta$. We counted the number of times this condition was satisfied in an auxiliary variable a . Finally, we returned $p_{TP} = \frac{a}{m}$.

The experiments were conducted by computing the p_{TP} of several different configurations. We precomputed 100,000 random samples. Then, we used a number of random seed document sets and ran the optimization routine for different configurations. We finally averaged the p values across the documents.

One of the configurations relates to the publication dates of the articles. The dates are scaled to a continuous range. In Fig. 8(a), it is evident that the p_{TP} value significantly decreases while increasing the *maximum date* value. This is because the range for the turning points increases with the scaling. Fig. 8(a) also shows that our approach generates statistically significant results with any value of the variance for the gamma function σ^2 , and that the change in significance while varying this parameter is limited.

The standard deviation parameter for overlap σ controls the width of the bell-shaped curve shown in Fig. 4. A penalty is added if any of the previous turning points fall within this curve. Figure 8(c) shows that as the overlap σ increases, the p_{TP} value increases. The maximum topical divergence parameter is used during the candidate generation stage to filter documents that are above this value. Figure 8(b) shows that when the topical divergence increases, the p_{TP} value decreases, albeit slightly. In general, our approach generates statistically significant results even with varying parameters.

6.5 Local Optimization: Repeatability of Concepts

One of the crucial elements of local optimization, as used in our framework, is that it may produce multiple results with different executions for the same set of seed documents and configurations. From an analytic perspective, multiple results for the same seed set provide a deeper understanding of the evolution. However, too much diversity in the results for the same seed set may be overwhelming for users. In this experiment, we analyze the repeatability of the results, in particular for the turning points, for the same set of seed sets. We selected random seed documents from different topics and repeated the optimization 100 times for each seed keeping track of the resulting turning-point vector, each time with a different initialization for the optimization routine. The collected vectors were compared pairwise and grouped into *buckets* based

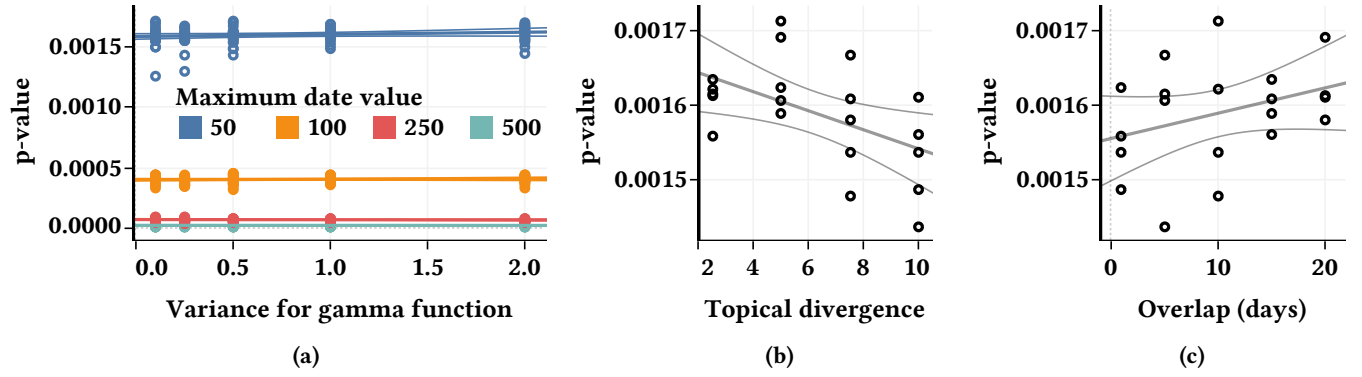


Figure 8: Significance of turning points vector vs. (a) variance of gamma function, (b) topical divergence and (c) date overlap penalty. Each data series includes a predicted trend line as well as upper and lower 95% confidence lines.

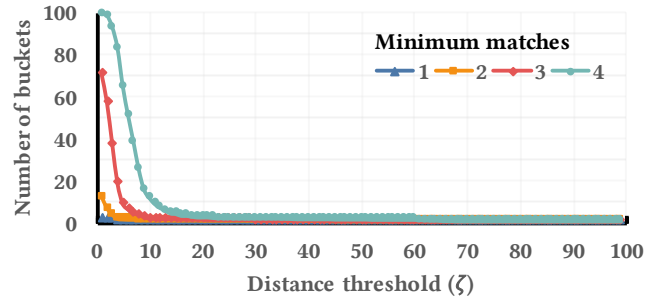


Figure 9: Number of buckets as a function of distance threshold ζ , with varying minimum matches.

on similarity. To define similarity in this context, we first define a match as when the distance between two turning points from different vectors is below a threshold, ζ . Two vectors are similar if the number of matches is above a *minimum match* parameter. Ideally, the number of buckets should be small (close to 1). If the number of buckets is 100, this means that none of the pairs of vectors fulfill the conditions mentioned above.

In Fig. 9, we observe how the average of the number of buckets for the seed documents changes with respect to the distance threshold ζ and a *minimum matches* score. ζ is varied from one to 100, since this is the range of the publication dates in this experiment, while the number of *minimum matches* is changed from one to four. Figure 9 shows that as the distance threshold increases, the number of buckets rapidly decreases. This indicates that, even with different initializations of the optimization routine with multiple executions for the same seed, we obtain similar chains and time segments. This helps the user to keep the analysis focused on a relevant set of evolving concepts for the same seed document.

6.6 Prediction

Our framework provides past documents as an evolution of the story of a given seed document. All the past documents have associated weights reflecting their importance in the evolution. These relevant documents can be used to study the evolution over time of the prominent entities found in the seed document. In this experiment, we examine if we can leverage these relevant documents to not only study the evolution of the entities in the past, but also to predict, using a simple linear regression, the future

evolution. A more specific question is: which entities will appear in an article published in the future given a seed document?

For this experiment, we divided the set of relevant documents based on the segments into a training dataset (four segments) and a testing dataset (one segment). Next, we built a table which consists of all pairs of words in all pairs of the documents of the training set. This table contains the difference in publication dates, as well as a pair of entities and corresponding tf-idf weights that appear in a pair of documents. That is, the training data reflects how far in time a pair of entities may appear and what are the tf-idf scores of those entities when they appear. We built a linear regression model for each of the terms appearing in the *future* documents, taking into account the difference in publication dates as an extra feature.

When a new prediction is requested, we use the linear regression models that include the entities of the seed document, to predict the weight of the entities in the *future* document. In Fig. 10(a), we show the word-cloud of a sample seed document. We then use our algorithm to predict how the word-clouds of three future documents would look like in a range between 4 and 10 days after the publication date of the original article. The actual and predicted word-clouds are presented in Fig. 10(b). As can be seen, even with a simple linear regression-based prediction model, for most of the cases the predicted documents have a match of at least two of the entities from the ground truth (e.g., *Paris attack*, *Belgium*, and *Paris*). The advantage of using the output of our model for prediction is that we exploit both the similarity of documents within a segment, and the diffusion of a topic across segments, to obtain better results.

We expect that a more sophisticated algorithm would be able to improve these results. The main disadvantage of this method is that if no entities from the seed document appear in the training set, then no entities can be predicted, and of course, unseen entities will not be predicted.

7 CONCLUSIONS AND FUTURE WORK

Our framework discovers the evolution of today's news stories from large news archives. We have presented case studies that demonstrate the possibility of using the resulting evolution-related documents as training data to predict the relevant entities for future news articles. Based on our case studies, we plan to expand our work in the direction of prediction where our framework will be able to generate warnings regarding insurgent activities, predict a



(a) A word-cloud of a seed document.



(b) On the left, we show the ground truth while on the right we show the predicted word-clouds. The words are sized according to the original and predicted weights, respectively. The publication date difference between the seed document and the predicted documents ranges between 4 and 10 days.

Figure 10: An experiment on predicting a future document.

timeline for developing a new concept, and forecast emerging stories. Furthermore, we plan to incorporate user feedback to improve the performance of our framework by modifying the optimization hyperparameters to meet the user's expectations.

REFERENCES

- [1] A. Ahmed, Q. Ho, J. Eisenstein, E. Xing, A. J. Smola, and C. H. Teo. 2011. Unified Analysis of Streaming News. In *WWW '11*. ACM, New York, NY, USA, 267–276.
- [2] Alias-i. 2008. LingPipe 4.1.0. (2008). <http://alias-i.com/lingpipe/>
- [3] J. Allan, R. Gupta, and V. Khandelwal. 2001. Temporal Summaries of New Topics. In *SIGIR '01*. ACM, New York, NY, USA, 10–18.
- [4] J. J. Angulo, C. A. Pederneiras, W. Ebner, E. M. Kimura, and P. Megale. 1980. Concepts of diffusion theory and a graphic approach to the description of the epidemic flow of contagious disease. *Public Health Reports* 95, 5 (1980), 478–485.
- [5] G. Binh Tran. 2013. Structured Summarization for News Events. In *WWW '13 Companion*. ACM, New York, NY, USA, 343–348.
- [6] D. M. Blei, A. Y. Ng, and M. I. Jordan. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* 3 (March 2003), 993–1022.
- [7] C. Faloutsos, K. S. McCurley, and A. Tomkins. 2004. Fast Discovery of Connection Subgraphs. In *KDD '04*. ACM, New York, NY, USA, 118–127.
- [8] L. Fang, A. D. Sarma, C. Yu, and P. Bohannon. 2011. REX: Explaining Relationships Between Entity Pairs. *Proc. VLDB Endow.* 5, 3 (Nov. 2011), 241–252.
- [9] J. R. Finkel, T. Grenager, and C. Manning. 2005. Incorporating Non-local Information into Information Extraction Systems by Gibbs Sampling. In *ACL '05*. ACL, Stroudsburg, PA, USA, 363–370.
- [10] J. Gillenwater, A. Kulesza, and B. Taskar. 2012. Discovering Diverse and Salient Threads in Document Collections. In *EMNLP-CoNLL '12*. ACL, Stroudsburg, PA, USA, 710–720.
- [11] W. Gu, S. Dong, and M. Chen. 2016. Personalized news recommendation based on articles chain building. *Neural Computing and Applications* 27, 5 (July 2016), 1263–1272.
- [12] K. Heath, N. Gelfand, M. Ovsjanikov, M. Aanjaneya, and L. J. Guibas. 2010. Image webs: Computing and exploiting connectivity in image collections. In *CVPR '10*. 3432–3439.
- [13] M. S. Hossain, C. Andrews, N. Ramakrishnan, and C. North. 2011. Helping Intelligence Analysts Make Connections. In *AAAIWS'11*. AAAI Press, Menlo Park, California, 22–31.
- [14] M. S. Hossain, P. Butler, A. P. Boedihardjo, and N. Ramakrishnan. 2012. Storytelling in Entity Networks to Support Intelligence Analysts. In *KDD '12*. ACM, New York, NY, USA, 1375–1383.
- [15] M. S. Hossain, J. Gresock, Y. Edmonds, R. Helm, M. Potts, and N. Ramakrishnan. 2012. Connecting the Dots between PubMed Abstracts. *PLOS ONE* 7, 1 (Jan. 2012), e29509.
- [16] Y. Jo, J. E. Hopcroft, and C. Lagoze. 2011. The Web of Topics: Discovering the Topology of Topic Evolution in a Corpus. In *WWW '11*. ACM, New York, NY, USA, 257–266.
- [17] D. Kim and A. Oh. 2011. Topic Chains for Understanding a News Corpus. In *CICLing '11*. Springer-Verlag, Berlin, Heidelberg, 163–176.
- [18] J. Kleinberg. 2002. Bursty and Hierarchical Structure in Streams. In *KDD '02*. ACM, New York, NY, USA, 91–101.
- [19] S. Kullback and R. A. Leibler. 1951. On Information and Sufficiency. *The Annals of Mathematical Statistics* 22, 1 (1951), 79–86.
- [20] D. Kumar, N. Ramakrishnan, R. F. Helm, and M. Potts. 2008. Algorithms for Storytelling. *IEEE Trans. Knowl. Data Eng.* 20, 6 (June 2008), 736–751.
- [21] Erdal Kuzey, Jilles Vreeken, and Gerhard Weikum. 2014. A Fresh Look on Knowledge Bases: Distilling Named Events from News. In *CIKM '14*. ACM, New York, NY, USA, 1689–1698.
- [22] X. Luo, J. Xuan, J. Lu, and G. Zhang. 2016. Measuring the Semantic Uncertainty of News Events for Evolution Potential Estimation. *ACM Trans. Inf. Syst.* 34, 4 (June 2016), 24:1–24:25.
- [23] M. A. Kader, S. M. Naim, A. P. Boedihardjo, and M. S. Hossain. 2016. Connecting the Dots Using Contextual Information Hidden in Text and Images. *AAAI Conference on Artificial Intelligence* (2016).
- [24] R. Nallapati, A. Feng, F. Peng, and J. Allan. 2004. Event Threading Within News Topics. In *CIKM '04*. ACM, New York, NY, USA, 446–453.
- [25] A. Nenkova and K. McKeown. 2012. A Survey of Text Summarization Techniques. In *Mining Text Data*, Charu C. Aggarwal and ChengXiang Zhai (Eds.). Springer US, 43–76.
- [26] Y. Ning, S. Muthiah, R. Tandon, and N. Ramakrishnan. 2015. Uncovering News-Twitter Reciprocity via Interaction Patterns. In *ASONAM '15*. ACM, New York, NY, USA, 1–8.
- [27] D. Radev, J. Otterbacher, A. Winkel, and S. Blair-Goldensohn. 2005. NewsIn-Essence: Summarizing Online News Topics. *Commun. ACM* 48, 10 (Oct. 2005), 95–98.
- [28] K. Radinsky, S. Davidovich, and S. Markovitch. 2012. Learning Causality for News Events Prediction. In *WWW '12*. ACM, New York, NY, USA, 909–918.
- [29] M. Rospocher, M. van Erp, P. Vossen, A. Fokkens, I. Aldabe, G. Rigau, A. Soroa, T. Ploeger, and T. Bogaard. 2016. Building Event-centric Knowledge Graphs from News. *Web Semant.* 37, C (March 2016), 132–151.
- [30] D. Shahaf and C. Guestrin. 2010. Connecting the Dots Between News Articles. In *KDD '10*. ACM, New York, NY, USA, 623–632.
- [31] D. Shahaf, C. Guestrin, E. Horvitz, and J. Leskovec. 2015. Information Cartography. *Commun. ACM* 58, 11 (Oct. 2015), 62–73.
- [32] C. Suen, S. Huang, C. Eksombatchai, R. Sosc, and J. Leskovec. 2013. NIFTY: A System for Large Scale Information Flow Tracking and Clustering. In *WWW '13*. ACM, New York, NY, USA, 1237–1248.
- [33] X. Wang, C. Zhai, and D. Roth. 2013. Understanding Evolution of Research Themes: A Probabilistic Generative Model for Citations. In *KDD '13*. ACM, New York, NY, USA, 1115–1123.
- [34] C. Wu, B. Wu, and B. Wang. 2016. Event Evolution Model Based on Random Walk Model with Hot Topic Extraction. 10086 (Dec. 2016), 591–603.
- [35] R. Yan, X. Wan, J. Otterbacher, L. Kong, X. Li, and Y. Zhang. 2011. Evolutionary Timeline Summarization: A Balanced Optimization Framework via Iterative Substitution. In *SIGIR '11*. ACM, New York, NY, USA, 745–754.
- [36] Y. Yang, T. Ault, T. Pierce, and C. W. Lattimer. 2000. Improving Text Categorization Methods for Event Tracking. In *SIGIR '00*. ACM, New York, NY, USA, 65–72.
- [37] Y. Yang, J. G. Carbonell, R. D. Brown, T. Pierce, B. T. Archibald, and X. Liu. 1999. Learning Approaches for Detecting and Tracking News Events. 14, 4 (July 1999), 32–43.
- [38] S. Yu, X. Li, X. Zhao, Z. Zhang, and F. Wu. 2015. Tracking News Article Evolution by Dense Subgraph Learning. *Neurocomput.* 168, C (Nov. 2015), 1076–1084.
- [39] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal. 1997. Algorithm 778: L-BFGS-B: Fortran Subroutines for Large-scale Bound-constrained Optimization. *ACM Trans. Math. Softw.* 23, 4 (Dec. 1997), 550–560.
- [40] X. Zhu and T. Oates. 2014. Finding Story Chains in Newswire Articles Using Random Walks. *Information Systems Frontiers* 16, 5 (Nov. 2014), 753–769.